

Addressing Appearance Change in Outdoor Robotics with Adversarial Domain Adaptation

Markus Wulfmeier¹, Alex Bewley¹ and Ingmar Posner¹

Abstract—Appearance changes due to weather and seasonal conditions represent a strong impediment to the robust implementation of machine learning systems in outdoor robotics. While the model is optimised for the training domain it will deliver degraded performance in application domains that underlie distributional shifts caused by these changes. Traditionally, this problem has been addressed via the collection of labelled data in multiple domains or by imposing priors on the type of shift between both domains. We frame the problem in the context of unsupervised domain adaptation and apply an adversarial framework to train a deep neural network with the additional objective to align features across domains. This approach benefits from adding unlabelled data and is generally applicable to many state-of-the-art architectures. Moreover, as adversarial training is notoriously hard to stabilise, we first perform an extensive ablation study on a surrogate classification task underlying the same appearance change and then apply the distilled insights to the problem of free-space segmentation for motion planning.

I. INTRODUCTION

In this paper we revisit the prominent issue in field robotics of appearance change under the influence of numerous factors including time of day, weather, and seasonal variation. Dealing with these changes becomes relevant in all modules of the robotic perception system including localisation [1], mapping [2] and obstacle detection [3]. Given the recent adoption of high capacity deep neural networks for many robotic vision tasks, the effects of condition variation have been moderately alleviated with the assumption that the labelled training data is diverse enough to capture the variation expected during deployment. However, due to the expense and impracticality of collecting labels across all environmental conditions, the training data is commonly captured in a subset of episodes which exacerbates the effects of appearance change when deployed.

This work addresses the problem of appearance change by reframing it in the context of unsupervised domain adaptation. With this view, we are acknowledging that the underlying distribution of our labelled training data may differ considerably from unlabelled data encountered in the application domain. Methods for domain adaptation have already found success in robotics for transfer from 3D models to laser data [4] or from simulation to real images collected indoors [5]. Here we focus on developing a general and flexible framework for adapting supervised machine learning models to address appearance change for outdoor robotics problems (as depicted in Figure 1).

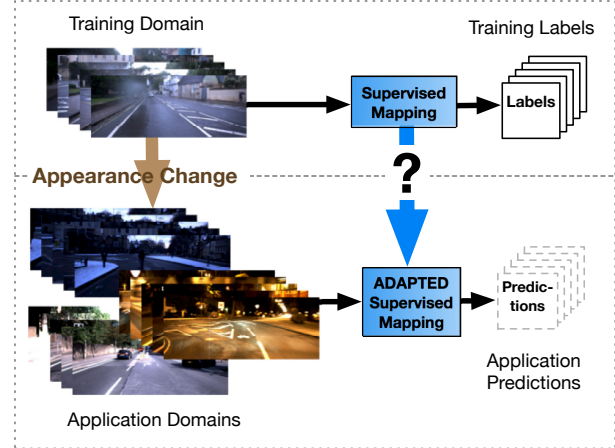


Fig. 1: Appearance change between training and test time of machine learning systems. Adversarial domain adaptation (ADA) provides a robust, systematic way to adapt the supervised task to perform in the application domains without supervised label information.

Traditional approaches for addressing the general domain adaptation problem have focused on modelling the density of the source and target distributions separately [6] or with imposed prior structure [7]. However, density estimation itself is a challenging issue, particularly for appearance, as it is difficult to impose a prior without making assumptions on the distribution. Recently, generative adversarial networks (GAN) were proposed [8] as a framework to model any arbitrary distribution where a generating network is optimised to produce data indistinguishable from real data as considered by a discriminating network. Building on the flexibility of GANs, adversarial domain adaptation (ADA) has demonstrated astonishing performance for unsupervised domain adaption [9], [10].

We develop a straightforward framework for the adaptation of existing, commonly used network architectures that allows to benefit from recent progress in deep learning. To demonstrate the widespread applicability of ADA with this schema, we modify two popular network architectures, AlexNet [11] and FCN-VGG16 [12] which are known to work well for the tasks of classification and pixel-wise image segmentation respectively. When considered in an outdoor context, both tasks are affected by change in appearance caused by factors such as time of day or weather conditions.

Due to the documented evidence of training instability of adversarial training procedures [8], [13], [14], we first use place classification as a surrogate task to perform an extensive evaluation of approaches for balancing our objectives and

¹The authors are with the Oxford Robotics Institute, Department of Engineering Science, University of Oxford, United Kingdom; markus, bewley, ingmar@robots.ox.ac.uk

stabilising the training process. The reduction in complexity enables us to enumerate and explore various potential configurations of the network structure and training procedures to identify the key factors of influence before scaling up to the more challenging segmentation task. Particularly, we seek to answer the following questions:

- 1) Trade-offs: How do we stabilise training and balance supervised and adversarial objectives?
- 2) Architecture: How can we directly adapt existing network architectures that are known to perform well?
- 3) Performance: Where does ADA work best and how do larger appearance changes influence performance?

A series of experiments is presented to provide valuable insights around the above questions facilitating the extension to full image pixelwise traversability segmentation commonly used as an input to motion planning systems. In the final segmentation task we see a significant performance increase in the target domain by applying ADA over the FCN-VGG16 [12] baseline. Furthermore, we also demonstrate that the approach can act as a regulariser and leads to performance gains within the original domain as more unlabelled data become available.

II. RELATED WORK

In robotics, the problem of appearance change has long been an issue for both laser and image data. Lai and Fox [4] presented an exemplar based approach to align spin-image features from web based 3D models to dense laser scan. For visual appearance, the colour-constant transform [15] is commonly applied directly on the image to achieve lighting invariance for localisation [16], visual teach and repeat [1] and segmentation [3]. Other works in localisation take a multi-view approach by either exploiting temporal structure [17] or by accumulating multiple experiences covering different appearances [2] to avoid the issue of modelling appearance change. The work of Neubert *et al.* [18] is the closest to our approach where they learn synthesise stored images into the current season. By casting the appearance change issue as an unsupervised domain adaption problem our approach differs from the previous work as we do not require known cross domain correspondences.

In the context of domain shift between training and test data, Ben-David *et al.* [19] derived theoretical upper bounds on a classifiers performance. While unsupervised domain adaptation is an open research problem in theoretical and practical terms [10], [20], recent successes have shown the capability to train expressive, flexible models and address high dimensional input distributions and therefore made first steps in enabling real world applicability [9], [21].

The majority of recent successes have built on the empirical superiority of Neural Networks that enable hierarchical representation learning with respect to arbitrary combinations of differential objectives. In the context of domain adaptation we can minimise the mutual information between feature representations and the underlying domain while maximising relevance towards a supervised objective in the source domain. Long et al [22] focus on minimising the Maximum Mean

Discrepancy for the feature distributions of multiple layers of the network architecture. Sun *et al.* [23] align second order statistics of layer activations for source and target domain.

Recently, the field has been extended with adversarial methods to domain adaptation [24], [25] which have resulted in strong performance benefits [10]. Adversarial unsupervised domain adaptation relies on training a discriminator to differentiate the domains underlying feature distributions while competing against an encoder which attempts to conceal the origin domain of input samples. While the loss formulation for the discriminator stays consistent between most approaches, different objectives have been applied for the encoder, including the *minimax* formulation resulting from a ‘gradient reversal layer’ [24], and the *confusion loss* [26] that was found to address problems with vanishing gradients based on discriminator saturation in Generative Adversarial Networks (GAN) [8]. These methods have recently been employed to address tasks in medical image segmentation [27], gaze estimation [9] and transfer for reinforcement learning [5].

III. ADVERSARIAL DOMAIN ADAPTATION

The principal goal of ADA is to maximise the performance of a supervised task not only in the source domain - where labels are available - but also in the unlabelled target domain. As a domain shift potentially exists between training and application domains, the approach tries to align the marginal feature distributions independent of label for both domains to achieve its goal.




Fig. 2: An illustration of domain shift where the distribution of the unlabelled target domain is misaligned with the labelled source domain but both distributions share structural similarities. By aligning the two domains the decision boundary of the supervised task (denoted in colour) better aligns across both domains.

Therefore, we train the supervised task module and encoder to maximise the likelihood of source labels given the source inputs. Additionally, to adapt the model towards performing well in the target domain, we train the encoder to confuse a discriminator which tries to estimate the domain of a data sample and serves as density model.

The adversarial training process works towards aligning the marginal feature distributions of source and target domain which given partial structural similarity of both domains implicitly also aligns the conditional distributions given the labels. Finally, by aligning the distributions, ADA increases the performance of a supervised module with decision boundaries optimised for the source domains as displayed in Figure 2.

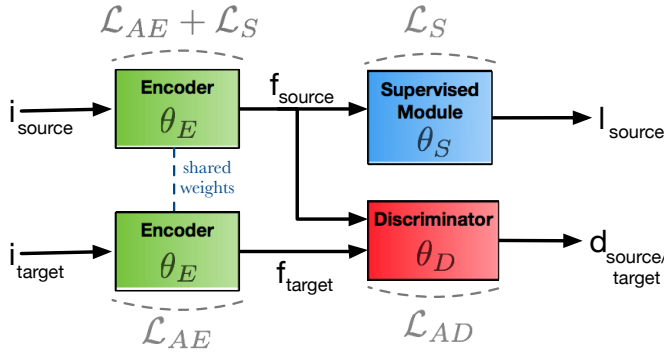


Fig. 3: Schema for adversarial domain adaptation. Given a well-established network architecture, we determine a split layer and pass the activations of that layer additionally to a discriminator module. While the supervised module is only applied to source data, both domains share the same encoder.

An existing network architecture is divided into two functional blocks referred to as *encoder* and *supervised task* as represented in Figure 3. Let $E : \mathbb{R}^n \rightarrow \mathbb{R}^m$ be the encoder that transforms an input image i_i into the feature representation f_i , which subsequently serves as input for both the supervised task $S : \mathbb{R}^m \rightarrow \mathbb{R}^c$ producing label l_i and the adversarial domain discriminator $D : \mathbb{R}^m \rightarrow \mathbb{R}$ computing domain label d_i . To reduce the memory footprint of the model, we can remove the discriminator at test time as it is only an auxiliary module to determine the encoder objective.

The model tries to simultaneously minimise supervised and adversarial losses, respectively \mathcal{L}_S and \mathcal{L}_A , where the adversarial loss is furthermore divided into the loss applying to the encoder \mathcal{L}_{AE} and discriminator \mathcal{L}_{AD} (see Equations 1 and 3). The supervised objective is minimising the cross-entropy loss from Equation 2. While the adversarial encoder and discriminator losses both depend on encoder parameters θ_E and discriminator parameters θ_D , each loss is only applied to the corresponding module to realise the adversarial training procedure. The filters of the supervised module θ_S are only optimised with respect to the supervised loss. The factor λ determines the relative strength of supervised and adversarial objective. In this arrangement, the encoder is encouraged to extract features that balance the relevance the supervised task and the maximisation of domain invariance.

$$\mathcal{L}(\theta_S, \theta_D, \theta_E) = \mathcal{L}_S(\theta_S, \theta_E) + \lambda \mathcal{L}_A(\theta_D, \theta_E) \quad (1)$$

$$\mathcal{L}_S(\theta_S, \theta_E) = \mathbb{E}_{l=S(E(i, \theta_E), \theta_S), i \sim S} [-\log(l)] \quad (2)$$

$$\mathcal{L}_A(\theta_D, \theta_E) = \mathcal{L}_{AD}(\theta_D) + \mathcal{L}_{AE}(\theta_E) \quad (3)$$

$$\begin{aligned} \mathcal{L}_{AD}(\theta_D) &= \mathbb{E}_{f=E(i), i \sim S} [-\log(D(f, \theta_D))] + \\ &\quad \mathbb{E}_{f=E(i), i \sim T} [-\log(1 - D(f, \theta_D))] \end{aligned} \quad (4)$$

Most existing work describes the adversarial encoder loss via a gradient reversal layer simply as the negated discriminator loss. To retain consistent with multiple possibilities known from the GAN framework, we treat it separately in our formulation.

The original generator loss in the GAN framework, similar to the resulting encoder loss in the ADA framework when

applied with gradient reversal, is the minimax loss in Equation 5. Contrary to the GAN framework, the loss applies to samples from both domains instead of only applying to generated samples. An alternative formulation is to train the encoder to maximise the discriminator’s confusion [8], [26] (see Equation 6). Both approaches are evaluated in Section IV-A.

$$\begin{aligned} \text{minimax loss} &: \\ \mathcal{L}_{AE}(\theta_E) &= -\mathcal{L}_{AD} \end{aligned} \quad (5)$$

$$\begin{aligned} &= \mathbb{E}_{f=E(i, \theta_E), i \sim S} [\log(D(f))] \\ &\quad + \mathbb{E}_{f=E(i, \theta_E), i \sim T} [\log(1 - D(f))] \end{aligned}$$

$$\begin{aligned} \text{confusion loss} &: \\ \mathcal{L}_{AE}(\theta_E) &= -\mathbb{E}_{f=E(i, \theta_E), i \sim S} [\log(1 - D(f))] \\ &\quad - \mathbb{E}_{f=E(i, \theta_E), i \sim T} [\log(D(f))] \end{aligned} \quad (6)$$

IV. ABLATION STUDY

Unsupervised domain adaptation is an inherently complex task as the lack of labelled information complicates the alignment between feature representations from source and target domain. While a high capacity encoder such as a deep neural network is per se capable of modelling even complex relations, the process of alignment might remove important information if it simplifies the alignment. The difficulty of the process is proportional to the difference between the underlying source and target domains. The field of autonomous driving presents a strong opportunity here as the use of overlapping routes can provide a partial alignment which simplifies the overall learning process.

We focus for our main evaluation on the surrogate task of classification and extend the evaluation subsequently to image segmentation with focus on path proposals for autonomous driving [28]. The hyperparameter study is built on a small subset of labelled data from the publicly available *Oxford RobotCar Dataset* that includes over 1000 km of driving data including images, LIDAR, GPS and INS data [29].

The principal experiments were performed based on 9,000 training and 1,000 test images for each domain and 20 distinct location classes. We focus on the adaptation between overcast weather to sunny as displayed in Figure 4.

The network architecture builds on AlexNet [11] and while we adapt the split between encoder and classifier / discriminator, the overall pipeline from image to location label is kept the same for all experiments, making this approach easy to apply to other common architectures such as FCN-VGG [12], which underlies the segmentation experiments.

We determine the mean classification accuracy P_T and standard deviation σ for all tested hyperparameter configurations over 5 runs to investigate for GAN-typical training instability. All experiments are performed on an NVIDIA GTX TITAN GPU. All evaluations included in this work are run with the best found set of hyperparameters for the fixed parameters in each test.

The following subsections now address the questions posed in Section I.



Fig. 4: Example images for source domain: overcast (left), and target domain: sunny (right).

A. Trade-offs: How do we stabilise training and balance supervised and adversarial objectives?

One of the most common issues with ADA is the potential inability to learn task-relevant and informative features as the domain confusion loss tries to reduce discriminator performance and can result in degenerate feature representations.

1) *Stabilising: Pretraining and Supervised Warm-up:* Initialising network architectures via pretraining of convolutional layers on large and diverse datasets is generally known to speed up the learning process as well as leading to better generalisation [30]. Furthermore, in the context of ADA, it can be helpful to include a warmup phase where only the supervised task loss is used before switching on the adversarial loss. We evaluate pretrained convolutional layers for AlexNet (based on ImageNet classification task [31]) as well as a supervised warmup phase of 15 epochs which, based on a small test run, has been found to work well in a large range of evaluation settings.

Table I shows that for the given small number of samples, supervised warmup on its own does not suffice for a strong, classifier relevant feature representation therefore we build our following evaluation on pretrained convolutional layers as well as a 15 epochs supervised warmup phase.




Fig. 5: Effect of different encoder loss formulations on the gradients. The confusion loss results in significantly stronger gradients against a well performing discriminator (close to 1 for source data and 0 for target data).

2) *Balancing: Adversarial Loss:* While the discriminator loss stays fixed across most recent work on ADA based on Equation 4, the choice of the encoder loss varies and has been shown to have significant influence on convergence properties and stability of training in Generative Adversarial Networks [14], which have strong similarities with ADA. The two main generator objectives to consider are the minimax formulation (negated discriminator loss) which is equal to

the gradient reversal layer [24] and the confusion loss [26], which has found to prevent vanishing gradients with saturated discriminator but displays higher variance in the gradients [14]. Additionally, the factor λ is used to vary the relative strength of the adversarial loss in the overall training process to balance both objectives.

Figure 5 shows minimax and confusion loss and gradients in dependence to the discriminator performance on source and target domain data. When the discriminator achieves high performance (close to 1 for source data and 0 for target data), the confusion loss results in significantly stronger gradients which will support a more stable adversarial training process.

	Supervised Warmup	Pretraining	Both
P_T [%]	52.68	72.03	82.03
σ_T [%]	6.57	4.88	2.09

TABLE I: Evaluation of performance for network initialisation and warmup phase. The combination of both approaches leads to significant improvement

Loss λ	Minimax Loss						
	10^{-2}	10^{-3}	10^{-4}	10^{-5}	10^{-6}	10^{-7}	10^{-8}
P_T [%]	59.68	55.22	64.35	73.28	73.66	76.32	68.76
σ_T [%]	2.05	3.40	4.73	3.76	1.89	2.13	3.62

Loss λ	Confusion Loss						
	10^1	10^0	10^{-1}	10^{-2}	10^{-3}	10^{-4}	10^{-5}
P_T [%]	76.04	82.03	71.89	71.95	71.28	53.93	37.77
σ_T [%]	1.14	2.09	2.71	2.00	1.64	29.06	29.71

TABLE II: Evaluation of performance for minimax and confusion loss with variation of loss weighting. Optimal performance is achieved with the confusion.

Layer	1	2	3	4	5	6	7
P_T [%]	72.25	74.14	78.93	82.03	79.37	79.72	78.63
σ_T [%]	1.75	1.71	2.90	2.09	1.07	0.98	0.84

TABLE III: Evaluation of performance for the variation of the split layer between encoder and classifier/discriminator. The optimal split lies in the middle of the network. While a split to early in the architecture can lead to low performance, moving the split further towards the last layers only slowly affects performance.

Layer Adapted Capacity	3			4			5		
	-2	0	2	-2	0	2	-2	0	2
P_T [%]	73.79	76.30	76.52	74.47	82.03	80.55	79.00	81.06	80.37
σ_T [%]	6.08	3.94	3.17	2.62	2.09	1.94	0.67	0.53	1.66

TABLE IV: Evaluation of performance for adapting the discriminator capacity. For the optimal split layer, the best performance is achieved by not separately changing the discriminator architecture.

The evaluation for the minimax loss in Table II begins with considerably lower weights as higher values for λ as 10^{-2}

resulted in significant training instability risking the collapse of the target performance to chance (5%). We have found the confusion loss significantly easier to apply as it renders the game between the adversarial modules more stable as displayed in Figure 5. The original minimax-loss formulation was successfully employed in a number of recent works [24], [10], however, we have found the tuning process more complex and obtained significantly stronger results based on the confusion loss.

B. Architecture: *How can we directly adapt existing network architectures that are known to perform well?*

When applying ADA to a new task, one can greatly benefit from applying existing network architectures that have been optimised and proven to perform well on similar tasks. This section addresses how to adapt existing architectures easily to incorporate the additional discriminator module.

The expressiveness and flexibility of encoder and discriminator directly influence the performance of our domain adaptation task. While the flexibility of the discriminator limits the types of domain discrepancies that can be detected, the encoder structure affects the efficiency of concealing the originating domain from the discriminator while generating relevant feature distributions for classification.

1) *Choice of Split Layer:* The following evaluation focuses in particular on AlexNet [11], which we adapt by providing the feature output of a particular layer additionally to a discriminator module. Following Figure 3, the encoder module now includes all layers before the split layer and is shared for source and target domain. We duplicate the architecture following the split layer for supervised module and discriminator and adapt the discriminator to output a single value per input as domain classifier. By varying the *split layer* with this approach we directly influence the capacity ratio between encoder and discriminator while keeping the overall number of layers for the supervised task fixed.

The best possible split results in the middle of the network as both - encoder and discriminator - have enough capacity to fulfil their respective tasks (see Table III). An earlier split layer significantly reduces the encoders capacity to minimise mutual information with respect to the domain while encoding classification-relevant features. A later split decreases the expressiveness of the discriminator’s density model such that less variations can be detected.

2) *Capacity of Discriminator:* While the choice of split layer as displayed above has strong influence on the effectiveness of domain adaptation, the approach might benefit from separately adapting the discriminator capacity to improve domain discrimination while keeping the rest of the architecture fixed. In this context, we now evaluate the performance with respect to varying capacity of the discriminator by either adding or removing 2 fully connected layers before the final layer.

Table IV presents that changing the capacity with the best split layer configuration (layer 4) leads to no improvement and it can be seen for all split layer configurations that trimming layers from the discriminator reduces target performance as



Fig. 6: Example images for source domain: day (left), and target domain: night (right).

the discriminator expressiveness is diminished. Furthermore, extending the discriminator leads to lowered accuracy as the model might overfit to domain discrepancy between both domain’s training sets.

C. Performance: *Where does ADA work best and how do larger appearance changes influence performance?*

To provide bounds for the performance of the approach we evaluate the classification accuracy of the classifier only trained on training data from the source domain as lower bound and only with training data from the target domain as upper bound. Note, all models are initialised with pretrained convolutional layers pretrained on ImageNet [31].

As the complexity for unsupervised domain adaptation directly correlates with the difference between the distributions, we evaluate benefits of the approach on the following source-target pairs *Sunny - Overcast* and *Day - Night* with the latter representing a significantly more complex transfer as visualised in Figure 6.

Classification				
	domains	AlexNet	AlexNet w ADA	AlexNet w target labels *
P_T [%]	overcast-sunny	67.95±1.02	82.03±2.09	(87.96±1.40)
P_T [%]	day-night	26.79±4.90	30.21±4.94	(90.42±1.04)
Free-Space Segmentation				
	FCV-VGG16	FCN- VGG16 w patch-ADA	FCN- VGG16 w ADA	FCN- VGG16 w target labels *
P_T [%]	75.12±0.76	67.68±4.25	85.27±1.03	(93.94±0.84)

TABLE V: Evaluation of the performance gains based on ADA for the surrogate classification task and free-space segmentation in the context of autonomous driving. ADA leads to significantly higher accuracy in all transfer domains and approaches in scenarios with minor shift towards the performance with known target labels. In the day-night scenario based on a strong difference in appearance, the approach only leads to minor benefits. The segmentation task measures performance P_T as mean average precision.

By increasing the domain shift occurring between source and target data, we can investigate limitations of the ADA. While the approach leads to performance gains in both scenarios, it is clearly better suited to address domains with limited

*Performance with available training labels on the target domain. This serves as upper bound for performance.

shift, such as the overcast-sunny scenario, and performs significantly closer to the upper bound obtained from training on labelled target data.

1) *Improving Performance in Source Domain:* To fully evaluate the suitability of ADA for long-term robotics applications, this section aims at investigating the source performance P_S as the original labelled domain is still of relevance. In addition to improving target performance, ADA can act as a regulariser to improve generalisation and test performance in the source domain as represented in Table VI. This gain in performance is only possible as long as both domains have significant structural similarities and the dominant variations are shared by both domains. With increasing discrepancy between the domains, it can however reduce performance on the source domain as it might diminish information that helps it to generalise in source but not in the target domain, as can be seen with respect to the day-night adaptation in Table VI.

	domains	trained on source	ADA
$P_S[\%]$	overcast-sunny	87.83 ± 1.56	90.82 ± 1.45
$P_S[\%]$	day-night	90.04 ± 3.01	89.16 ± 1.78

TABLE VI: Evaluation of the regularising effect of ADA on source performance. For minor appearance difference between the domains, the approach can improve source performance by acting as a regulariser. In case of stronger shifts such as the day-night transfer it can lead to reduced source performance though.

V. SEGMENTATION TASK

Following the optimisation for our surrogate task, we now apply the distilled insights for optimising ADA to the task of free space segmentation as possible input data for motion planning systems. We use the fully convolutional FCN-VGG16 [12] architecture, which is split - similar to the classification tests - into encoder and classifier/discriminator. We set the split layer towards the middle of the architecture after the 4th maxpool operation (see [12] for the exact architecture) with fixed capacity of the discriminator architecture and apply the confusion loss.

Both source and target datasets include 1000 training and 100 test images based on a midday to early evening adaptation scenario such that the domain shift is intuitively smaller than in the full day to night transfer from section IV-C. The segmentation labels are generated for free-space/obstacles following the approach of Barnes et al [28].

The domain adapted segmentation network performs significantly better than the basic supervised approach and is able to bridge the gap towards performance with available labels in the target domain.

As the segmentation output of the approach has a limited receptive field for each pixel location we additionally evaluate a patch-wise discriminator in line with research on image-to-image translation with conditional Generative Adversarial Networks (cGAN) [32]. This approach enables to keep the size of the receptive field fixed between the supervised and adversarial task. While Isola et al [32] have found the patch

discriminator to work better for their task based on cGANs, it resulted in reduced accuracy for our application on ADA.




Fig. 7: Examples of segmentation results in the target domain. Green and red represent free-space and obstacles respectively. ADA leads to significant performance gains for the unsupervised target domain while patch-ADA fails to increase accuracy. The scenario with available target labels serves as upper bound for the performance when supervised information would be available.

VI. LIMITATIONS OF UNSUPERVISED ADVERSARIAL DOMAIN ADAPTATION

While we have shown domain adaptation to be beneficial in many scenarios, it must be marked that the current approach finds its limitations when the differences between source and target domains are too severe. As exemplified by the day to night transfer scenario, the approach still leads to some improvement but with stronger variation in the underlying domains, the adversarial encoder loss might even lead to reduced performance if weighted improperly against the supervised loss. If the adversarial loss dominates in such situations the encoder features might lose more information relevant for the supervised task.

Finally, to overcome the limitation of significant domain shifts, semi-supervised approaches can be employed to incorporate further structure and align the conditional distributions over feature representations given the labels comparable to [26].

VII. DISCUSSION

Adversarial training frameworks such as ADA tend to be notoriously hard to train. However, a limited number of hyperparameters has strong influence for a given problem and can be adapted to stabilise and optimise the training process.

While the detailed performance depends on architecture and task, we found the principal factors for optimising performance to be:

- Using relevant initialisation including supervised warmup helps to guide the training process.
- Applying the confusion loss for the encoder enables better balancing and stabilising.
- The optimum position for the split layer is mid network. Particularly the application in earlier layers can significantly reduce the benefits.

Additional to the main evaluation above, we tested the influence of other advancements from the related GAN framework. It was found that neither mini-batch discriminator [33] nor discriminator noise [33] brought significant advances. This is justified as the former addresses generator mode collapse which in the context of ADA will be less of a problem as the supervised loss guides towards more versatile solutions. The latter seems to have negligible influence in comparison to the GAN framework as the discriminator's task is of higher complexity and with reasonable loss settings, the risk of the discriminator saturating is minimal.

As a side note, we find it helpful in this adversarial training framework to apply gradient clipping to prevent abrupt instabilities when one of the adversarial models finds a strong exploit.

VIII. CONCLUSION AND FUTURE WORK

In this paper, we cast the common problem of appearance change in outdoor robotics as a domain adaptation problem and employ the recent adversarial paradigm to enhance the performance of existing architectures in unlabelled domains. While instabilities of adversarial training can inhibit the extension to large scale problems, our extensive tests on a surrogate task with moderate complexity expose the most significant factors of influence. With this evaluation we hope to pave the way for further application on real world tasks, in particular in the context of autonomous mobility where strong structural similarities can exist between different domains, e.g. based on spatial overlay of driven routes.

Beyond dealing with appearance change we see many potential applications of ADA in robotics where sensor modalities may change, or even transferring models from a simulated virtual environment to improve their performance in the real world.

Finally, the application of ADA represents a strong opportunity to address curriculum learning to overcome strong domain shifts as were shown in this paper to produce significant challenges for direct adaptation. As many appearance changes happen continuously over time, one can imagine adapting models more gradually for changing daytime, season, or weather.

REFERENCES

- [1] Michael Paton, Kirk MacTavish, Chris J Ostafew, and Timothy D Barfoot. It's not easy seeing green: Lighting-resistant stereo visual teach & repeat using color-constant images. In *Robotics and Automation (ICRA), 2015 IEEE International Conference on*, pages 1519–1526. IEEE, 2015.
- [2] Winston Churchill and Paul Newman. Experience-based Navigation for Long-term Localisation. *The International Journal of Robotics Research (IJRR)*, 2013.
- [3] Will Maddern, Alex Stewart, Colin McManus, Ben Upcroft, Winston Churchill, and Paul Newman. Illumination invariant imaging: Applications in robust vision-based localisation, mapping and classification for autonomous vehicles. In *Proceedings of the Visual Place Recognition in Changing Environments Workshop, IEEE International Conference on Robotics and Automation (ICRA), Hong Kong, China*, volume 2, page 3, 2014.
- [4] Kevin Lai and Dieter Fox. 3d laser scan classification using web data and domain adaptation. In *Robotics: Science and Systems*, volume 2, 2009.
- [5] Eric Tzeng, Coline Devin, Judy Hoffman, Chelsea Finn, Xingchao Peng, Sergey Levine, Kate Saenko, and Trevor Darrell. Towards adapting deep visuomotor representations from simulated to real environments. *arXiv preprint arXiv:1511.07111*, 2015.
- [6] Hal Daume III and Daniel Marcu. Domain adaptation for statistical classifiers. *Journal of Artificial Intelligence Research*, 26:101–126, 2006.
- [7] Masashi Sugiyama, Shinichi Nakajima, Hisashi Kashima, Paul V Buenau, and Motoaki Kawanabe. Direct importance estimation with model selection and its application to covariate shift adaptation. In *Advances in neural information processing systems*, pages 1433–1440, 2008.
- [8] Ian Goodfellow, Jean Pouget-Abadie, Mehdi Mirza, Bing Xu, David Warde-Farley, Sherjil Ozair, Aaron Courville, and Yoshua Bengio. Generative adversarial nets. In *Advances in neural information processing systems*, pages 2672–2680, 2014.
- [9] Ashish Shrivastava, Tomas Pfister, Oncel Tuzel, Josh Susskind, Wenda Wang, and Russ Webb. Learning from simulated and unsupervised images through adversarial training. *CoRR*, abs/1612.07828, 2016.
- [10] Konstantinos Bousmalis, George Trigeorgis, Nathan Silberman, Dilip Krishnan, and Dumitru Erhan. Domain separation networks. In *Advances in Neural Information Processing Systems*, pages 343–351, 2016.
- [11] Alex Krizhevsky, Ilya Sutskever, and Geoffrey E Hinton. Imagenet classification with deep convolutional neural networks. In *Advances in neural information processing systems*, pages 1097–1105, 2012.
- [12] Jonathan Long, Evan Shelhamer, and Trevor Darrell. Fully convolutional networks for semantic segmentation. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 3431–3440, 2015.
- [13] Ian J. Goodfellow. NIPS 2016 tutorial: Generative adversarial networks. *CoRR*, abs/1701.00160, 2017.
- [14] Martin Arjovsky and Léon Bottou. Towards principled methods for training generative adversarial networks. In *NIPS 2016 Workshop on Adversarial Training. In review for ICLR*, volume 2016, 2017.
- [15] Sivalogeswaran Ratnasingham and Steve Collins. Study of the photodetector characteristics of a camera for color constancy in natural scenes. *JOSA A*, 27(2):286–294, 2010.
- [16] Peter Corke, Rohan Paul, Winston Churchill, and Paul Newman. Dealing with shadows: Capturing intrinsic scene appearance for image-based outdoor localisation. In *Intelligent Robots and Systems (IROS), 2013 IEEE/RSJ International Conference on*, pages 2085–2092. IEEE, 2013.
- [17] Michael J Milford and Gordon F Wyeth. Seqslam: Visual route-based navigation for sunny summer days and stormy winter nights. In *Robotics and Automation (ICRA), 2012 IEEE International Conference on*, pages 1643–1649. IEEE, 2012.
- [18] Peer Neubert, Niko Sunderhauf, and Peter Protzel. Appearance change prediction for long-term navigation across seasons. In *Mobile Robots (ECMR), 2013 European Conference on*, pages 198–203. IEEE, 2013.
- [19] Shai Ben-David, John Blitzer, Koby Crammer, Fernando Pereira, et al. Analysis of representations for domain adaptation. *Advances in neural information processing systems*, 19:137, 2007.
- [20] Konstantinos Bousmalis, Nathan Silberman, David Dohan, Dumitru Erhan, and Dilip Krishnan. Unsupervised pixel-level domain adaptation

with generative adversarial networks. *arXiv preprint arXiv:1612.05424*, 2016.

- [21] Sungeun Hong, Woobin Im, Jongbin Ryu, and Hyun S. Yang. SSPP-DAN: Deep Domain Adaptation Network for Face Recognition with Single Sample Per Person. page 5, feb 2017.
- [22] Mingsheng Long, Jianmin Wang, and Michael I. Jordan. Deep transfer learning with joint adaptation networks. *CoRR*, abs/1605.06636, 2016.
- [23] Baochen Sun, Jiashi Feng, and Kate Saenko. Correlation alignment for unsupervised domain adaptation. *CoRR*, abs/1612.01939, 2016.
- [24] Yaroslav Ganin, Evgeniya Ustinova, Hana Ajakan, Pascal Germain, Hugo Larochelle, François Laviolette, Mario Marchand, Victor Lempitsky, Urun Dogan, Marius Kloft, Francesco Orabona, and Tatiana Tommasi. Domain-Adversarial Training of Neural Networks. *Journal of Machine Learning Research*, 17:1–35, 2016.
- [25] Hana Ajakan, Pascal Germain, Hugo Larochelle, François Laviolette, and Mario Marchand. Domain-adversarial neural networks. *arXiv preprint arXiv:1412.4446*, 2014.
- [26] Eric Tzeng, Judy Hoffman, Trevor Darrell, and Kate Saenko. Simultaneous deep transfer across domains and tasks. In *Proceedings of the IEEE International Conference on Computer Vision*, pages 4068–4076, 2015.
- [27] Konstantinos Kamnitsas, Christian F. Baumgartner, Christian Ledig, Virginia F. J. Newcombe, Joanna P. Simpson, Andrew D. Kane, David K. Menon, Aditya Nori, Antonio Criminisi, Daniel Rueckert, and Ben Glocker. Unsupervised domain adaptation in brain lesion segmentation with adversarial networks. *CoRR*, abs/1612.08894, 2016.
- [28] Dan Barnes, William P. Maddern, and Ingmar Posner. Find your own way: Weakly-supervised segmentation of path proposals for urban autonomy. *CoRR*, abs/1610.01238, 2016.
- [29] Will Maddern, Geoff Pascoe, Chris Linegar, and Paul Newman. 1 Year, 1000km: The Oxford RobotCar Dataset. *The International Journal of Robotics Research (IJRR)*, 2016.
- [30] Mi-Young Huh, Pulkit Agrawal, and Alexei A. Efros. What makes imagenet good for transfer learning? *CoRR*, abs/1608.08614, 2016.
- [31] Olga Russakovsky, Jia Deng, Hao Su, Jonathan Krause, Sanjeev Satheesh, Sean Ma, Zhiheng Huang, Andrej Karpathy, Aditya Khosla, Michael Bernstein, et al. Imagenet large scale visual recognition challenge. *International Journal of Computer Vision*, 115(3):211–252, 2015.
- [32] Phillip Isola, Jun-Yan Zhu, Tinghui Zhou, and Alexei A Efros. Image-to-image translation with conditional adversarial networks. *arXiv preprint arXiv:1611.07004*, 2016.
- [33] Tim Salimans, Ian Goodfellow, Wojciech Zaremba, Vicki Cheung, Alec Radford, and Xi Chen. Improved techniques for training gans. In *Advances in Neural Information Processing Systems*, pages 2226–2234, 2016.